

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant : Thomas Michael Gil et al. Art Unit : 2155
Serial No. : 09/931,223 Examiner : Nawar, Asad M.
Filed : August 16, 2001 Conf. No. : 2855
Title : STATISTICS COLLECTION FOR NETWORK TRAFFIC

MAIL STOP APPEAL BRIEF – PATENTS

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

APPEAL BRIEF ON BEHALF OF THOMAS MICHAEL GIL ET AL.

The Appeal Brief fee of **\$250** has already been paid. If an additional fee is due please apply that Appeal Brief fee and any other charges or credits to Deposit Account No. 06-1050.

(i.) Real Party In Interest

The real party in interest in the above application is Mazu Networks, Inc.

(ii.) Related Appeals and Interferences

The appellant is not aware of any appeals or interferences related to the above-identified patent application.

(iii.) Status of Claims

This is an appeal from the decision of the Primary Examiner in an Office Action dated June 6, 2006, rejecting claims 1-21 and 50-77, all of the claims in the application. Claims 22-49 were canceled. The claims have been twice rejected. Claims 1-21 and 50-77 are the subject of this appeal.

(iv.) Status of Amendments

This Appeal Brief is accompanied by a Reply to correct minor typographical errors in claims 63 and 70. Since this amendment is being filed as a matter of right, the amendment will necessarily have to be entered by the examiner. Accordingly, all amendments have been entered. Appellant has filed herewith a new Notice of Appeal. Appellant previously filed a Notice of Appeal on **December 14, 2005** and an Appeal Brief on March 7, 2006. The examiner replied with the above identified office action from which Appellant now appeals from.

(v.) Summary of Claimed Subject Matter

Claim 1

One aspect of Appellant's invention is set out in claim 1 as a machine implemented method of monitoring traffic flow in a monitoring device disposed to receive network traffic packets. "Referring to FIG. 6, a monitoring process 32 is shown. The monitoring process 32 can be deployed on data collectors 28 as well as gateways 26." [Appellant's specification Page 13, lines 24-27].

Inventive features of claim 1 include producing statistics corresponding to a parameter of traffic flow to trace the source of an attack. "Referring to FIG. 4, the data collector 26 performs a sampling and statistic collection process 40. The data collector samples one (1) packet in every (n) packets and has counters to collect statistics about every packet." [Appellant's specification Page 9, lines 11-14]. "The gateways 26 and data collectors have monitoring process 32 used to measure some parameter of traffic flow. One goal of the gateways 26 and data collectors 28 is to measure some parameter of network traffic. This information collected by the gateways 26 and data collectors is used to trace the source of an attack." [Appellant's specification Page 14, lines 5-10].

Inventive features of claim 1 include mapping the traffic flow into a plurality of buckets by applying a hash function " $f(h)$ " to the parameter of the traffic flow to output an integer corresponding to one of the buckets. "The algorithm will use some hash function " $f(h)$ ", which takes the packet and outputs an integer that corresponds to one of the buckets " $B_1 - B_N$." [Appellant's specification Page 14, lines 18-21].

Inventive features of claim 1 include accumulating statistics from the packets and comparing the number of buckets to a threshold. "Statistics from the packets start accumulating in the buckets " $B_1 - B_N$ ". The buckets " $B_1 - B_N$ " are configured with threshold values "Th." As the contents of the buckets $B_1 - B_N$ reach the configured thresholds values "Th", (e.g., compare values of packet count or packet rate to threshold), the monitoring process 32 deems that event to be of significance." [Appellant's specification Page 14, lines 21-25].

Inventive features of claim 1 include determining whether the number of buckets should be divided into more buckets or combined into fewer buckets based on comparing the number of buckets to the threshold. "As the gateway 26 or data collector 28 approaches a bucket threshold "Th", the gateway 26 or data collector 28 have the ability to take several buckets $B_1 - B_3$ and divide them in more buckets $B_1 - B_4$ or combine them into fewer bucket $B_1 - B_2$." [Appellant's specification Page 15, lines 18-22].

Claim 14

Claim 14 claims another aspect of the invention. Claim 14 is a computer program product residing on a computer readable for monitoring network traffic flow in a network. [Appellant's specification Page 2, lines 8-12]. The gateway 26 and data collector 26 are typically software programs that are executed on devices such as computers, routers, or switches. [Appellant's specification Page 9, lines 6-8].

Inventive features of claim 14 include instructions to map traffic flow into a plurality of buckets by applying a hash function " $f(h)$ " to a parameter of the traffic flow to output an integer corresponding to one of the buckets. This feature is supported as the analogous feature of claim 1.

Inventive features of claim 14 include instructions to accumulate statistics from the packets and compare the accumulated statistic values from the buckets to configured threshold values corresponding to the number of buckets to determine that an event is of significance. This feature is supported as the analogous feature of claim 1.

Inventive features of claim 14 include instructions to adjust the number of buckets as the number of buckets approaches a second threshold. This feature is supported as the analogous feature of claim 1.

Claim 21

Another aspect of the invention is covered by claim 21. Claim 21 is directed to a data collector to collect statistical information about network flows. "Referring to FIG. 4, the data collector 26 performs 40 a sampling and statistic collection process 40. The data collector samples 42 one (1) packet in every (n) packets and has counters to collect statistics about every packet." [Appellant's specification Page 9, lines 11-14].

Inventive features of claim 21 include a computer readable medium and a computing device that executes a computer program product stored on the computer readable medium. "The gateway 26 and data collector 26 are typically software programs that are executed on devices such as computers, routers, or switches." [Appellant's specification Page 9, lines 6-8].

Inventive features of claim 21 include instructions to map traffic flow into a plurality of buckets by applying a hash function " $f(h)$ " to the parameter of the traffic flow to output an

integer corresponding to one of the buckets. This feature is supported as the analogous feature of claim 1.

Inventive features of claim 21 include instructions to accumulate statistics from the packets and compare the accumulated statistic values from the buckets to configured threshold values corresponding to the number of buckets to determine that an event is of significance, adjust the number of buckets as the number of buckets approaches a second threshold. This feature is supported as the analogous feature of claim 1.

Claim 63

Claim 63 is directed to a method of monitoring traffic flow in a monitor device disposed to receive network packets. This feature is supported as the analogous feature of claim 1.

Inventive features of claim 63 include producing statistics corresponding to a parameter of the traffic flow to trace a source of an attack. This feature is supported as the analogous feature of claim 1.

Inventive features of claim 63 include mapping the traffic flow into a plurality of buckets. This feature is supported as the analogous feature of claim 1.

Inventive features of claim 63 include varying the number of buckets according to the amount of traffic and number of flows to breakdown traffic flow into different buckets. "As the gateway 26 or data collector 28 approaches a bucket threshold "Th", the gateway 26 or data collector 28 have the ability to take several buckets $B_1 \dots B_3$ and divide them in more buckets $B_1 \dots B_4$ or combine them into fewer bucket $B_1 \dots B_2$. [Appellant's specification Page 15, lines 18-22].

Inventive features of claim 63 also include analyzing statistics accumulated for a parameter and a corresponding threshold in the bucket to identify the source of the attack. "The function of the variable number of buckets is to dynamically adjust the monitoring process to the amount of traffic and number of flows, so that the monitoring device (e.g., gateway 26 or data collector 28) is not vulnerable to DoS attacks against its own resources. The variable number of buckets also efficiently identifies the source(s) of attack by breaking down traffic into different

categories (buckets) and looking at the appropriate parameters and thresholds in each bucket.”
[Appellant's specification Page 15, lines 23-31].

Claim 70

Claim 70 is directed to a computer program product residing on a computer readable medium for monitoring traffic flow in a monitor device disposed to receive network packets. This feature is supported as the analogous feature of claim 1.

Inventive features of claim 70 include instructions to produce statistics corresponding to a parameter of the traffic flow to trace a source of an attack. This feature is supported as the analogous feature of claim 1.

Inventive features of claim 70 in addition include instructions to map the traffic flow into a plurality of buckets. This feature is supported as the analogous feature of claim 1.

Inventive features of claim 70 also include instructions to vary the number of buckets according to the amount of traffic and number of flows to breakdown the traffic flow into different buckets. This feature is supported as the analogous feature of claim 63.

Inventive features of claim 70 also include instructions to analyze statistics accumulated for a parameter and a corresponding threshold in the bucket to identify a source of the attack. This feature is supported as the analogous feature of claim 63.

(vi.) Grounds of Rejection to be Reviewed on Appeal

1. Claim 63 stands rejected under 35 U.S.C. 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention. More specifically it is not clear what is further comprising.

2. Claims 63-68 and 70-75 stand rejected under 35 U.S.C. 102(e) as being anticipated by Lyle et al (US 6,971,028) hereinafter referred to as Lyle.

3. Claims 1-21, 50-62, 69, and 76-77 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Lyle further in view of Hsu et al (US 6,098,157) hereinafter referred to as Hsu.

(vii.) **Argument**

Anticipation

"It is well settled that anticipation under 35 U.S.C. §102 requires the presence in a single reference of all of the elements of a claimed invention." *Ex parte Chopra*, 229 U.S.P.Q. 230, 231 (BPA&I 1985) and cases cited.

"Anticipation requires the presence in a single prior art disclosure of all elements of a claimed invention arranged as in the claim." *Connell v. Sears, Roebuck & Co.*, 220 U.S.P.Q. 193, 198 (Fed. Cir. 1983).

"This court has repeatedly stated that the defense of lack of novelty (i.e., 'anticipation') can only be established by a single prior art reference which discloses each and every element of the claimed invention." *Structural Rubber Prod. Co. v. Park Rubber Co.*, 223 U.S.P.Q. 1264, 1270 (Fed. Cir. 1984), citing five prior Federal Circuit decisions since 1983 including *Connell*.

In a later analogous case the Court of Appeals for the Federal Circuit again applied this rule in reversing a denial of a motion for judgment n.o.v. after a jury finding that claims were anticipated. *Jamesbury Corp. v. Litton Industrial Prod., Inc.*, 225 U.S.P.Q. 253 (Fed. Cir. 1985).

After quoting from *Connell*, "Anticipation requires the presence in a single prior art disclosure of all elements of a claimed invention arranged as in the claim," 225 U.S.P.Q. at 256, the court observed that the patentee accomplished a constant tight contact in a ball valve by a lip on the seal or ring which interferes with the placement of the ball. The lip protruded into the area where the ball will be placed and was thus deflected after the ball was assembled into the valve. Because of this constant pressure, the patented valve was described as providing a particularly good seal when regulating a low pressure stream. The court quoted with approval from a 1967 Court of Claims decision adopting the opinion of then Commissioner and later Judge Donald E. Lane:

[T]he term "engaging the ball" recited in claims 7 and 8 means that the lip contacts the ball with sufficient force to provide a fluid tight seal ***** The Saunders flange or lip only sealingly engages the ball 1 on the upstream side when the fluid pressure forces the lip against the ball and never sealingly engages the ball on the downstream side because

there is no fluid pressure there to force the lip against the ball. The Saunders sealing ring provides a compression type of seal which depends upon the ball pressing into the material of the ring. *** The seal of Saunders depends primarily on the contact between the ball and the body of the sealing ring, and the flange or lip sealingly contacts the ball on the upstream side when the fluid pressure increases. 225 U.S.P.Q. at 258.

Relying on *Jamesbury*, the ITC said, "Anticipation requires looking at a reference, and comparing the disclosure of the reference with the claims of the patent in suit. A claimed device is anticipated if a single prior art reference discloses all the elements of the claimed invention as arranged in the claim." *In re Certain Floppy Disk Drives and Components Thereof*, 227 U.S.P.Q. 982, 985 (U.S. ITC 1985).

Obviousness

"It is well established that the burden is on the PTO to establish a prima facie showing of obviousness, *In re Fritsch*, 972 F.2d 1260, 23 U.S.P.Q.2d 1780 (C.C.P.A., 1972)."

"It is well established that there must be some logical reason apparent from the evidence or record to justify combination or modification of references. *In re Regal*, 526 F.2d 1399 188, U.S.P.Q.2d 136 (C.C.P.A. 1975). In addition, even if all of the elements of claims are disclosed in various prior art references, the claimed invention taken as a whole cannot be said to be obvious without some reason given in the prior art why one of ordinary skill in the art would have been prompted to combine the teachings of the references to arrive at the claimed invention. *Id.* Even if the cited references show the various elements suggested by the Examiner in order to support a conclusion that it would have been obvious to combine the cited references, the references must either expressly or impliedly suggest the claimed combination or the Examiner must present a convincing line of reasoning as to why one skilled in the art would have found the claimed invention obvious in light of the teachings of the references. *Ex Parte Clapp*, 227 U.S.P.Q.2d 972, 973 (Board. Pat. App. & Inf. 985)."

"The mere fact that the prior art could be so modified would not have made the modification obvious unless the prior art suggested the desirability of the modification." *In re Gordon*, 221 U.S.P.Q. 1125, 1127 (Fed. Cir. 1984).

Although the Commissioner suggests that [the structure in the primary prior art reference] could readily be modified to form the [claimed] structure, "[t]he mere fact that the prior art could be so modified would not have made the modification obvious unless the prior art suggested the desirability of the modification." *In re Laskowski*, 10 U.S.P.Q. 2d 1397, 1398 (Fed. Cir. 1989).

"The claimed invention must be considered as a whole, and the question is whether there is something in the prior art as a whole to suggest the desirability, and thus the obviousness, of making the combination." *Lindemann Maschinenfabrik GMBH v. American Hoist & Derrick*, 221 U.S.P.Q. 481, 488 (Fed. Cir. 1984).

Obviousness cannot be established by combining the teachings of the prior art to produce the claimed invention, absent some teaching or suggestion supporting the combination. Under Section 103, teachings of references can be combined only if there is some suggestion or incentive to do so. *ACS Hospital Systems, Inc. v. Montefiore Hospital*, 221 U.S.P.Q. 929, 933 (Fed. Cir. 1984) (emphasis in original, footnotes omitted).

"The critical inquiry is whether 'there is something in the prior art as a whole to suggest the desirability, and thus the obviousness, of making the combination.'" *Fromson v. Advance Offset Plate, Inc.*, 225 U.S.P.Q. 26, 31 (Fed. Cir. 1985).

1. Claim 63, as amended, is proper under 35 U.S.C. 112, second paragraph.

Appellant has amended claims 63 and 70 in the accompanying Reply to clarify a phrase in the claims and delete extraneous words. The feature now recites: "vary the number of buckets according to the amount of traffic and number of flows to breakdown the traffic flow into different buckets.", e.g., for claim 63. As amended, claim 63 is proper under 35 U.S.C. 112,

second paragraph. Analogous amendments were made to claim 70, since claim 70 had similar typographical errors.

**2. Claims 63-68 and 70-75 are not
anticipated by Lyle et al (US 6,971,028).**

Claims 63, 66, 70 and 73

For the purposes of this appeal only claims 63, 66, 70 and 73 stand or fall together. Claim 63 is representative of this group of claims.

Claim 63 is directed to a method of monitoring traffic flow in a monitor device disposed to receive network traffic packets. Claim 63 includes the features of producing statistics corresponding to a parameter of traffic flow to trace the source of an attack. According to claim 63 producing includes mapping the traffic flow into a plurality of buckets and varying the number of buckets according to the amount of traffic and number of flows by breaking down traffic flow into different buckets and examining statistics accumulated for a parameter and a corresponding threshold in the bucket.

The examiner contends that Lyle teaches "Producing statistics corresponding to a parameter of traffic flow to trace the source of an attack (Fig 9, 908-310; col 2, lines 45-50; col 7, lines 3-12; sniffers are used in analyzing and evaluating traffic flows to scrutinize suspicious activity in an attempt to ascertain the source of an attack)"

Appellant disagrees. Lyle neither describes nor suggests producing statistics corresponding to a parameter of traffic flow. Lyle merely uses sniffers, but according to Lyle, the sniffer "continuously scans the data being received at various ports of various network devices. The sniffers search for data indicating an actual or suspected attack, as described more fully below, and provide information concerning suspicious data to other modules within the tracking system, as described more fully below." [Lyle Col. 7, Lines 7-12].

Sniffers in Lyle are used to examine data in packets that have the characteristics of a known attack. Lyle does not disclose the sniffers as collecting statistical information on network traffic seen at nodes.

The examiner argues that Lyle discloses: "Mapping the traffic flow into a plurality of buckets (col 7, lines 43-67; event data, which is defined as suspicious data is placed in a queue as a set corresponding to a single incident)." Appellant contends that Lyle does not disclose this feature either at Col .7, lines 43-67 or by the definition of event data: "defined as suspicious data is placed in a queue as a set corresponding to a single incident.", since claim 63 requires mapping the traffic flow into a plurality of buckets, not events that correspond to incidents. The events are not traffic flow.

The examiner argues that Lyle discloses: "Varying the number of buckets according to the amount of traffic and number of flows according to down traffic flow into different buckets and examining statistics accumulated for a parameter and a corresponding threshold in the bucket (col 7, line 43 to col 8, line 5; col 13, lines 42-50" once an event (a set of data corresponding to an attack) is placed in the queue, other event data is grouped or combined with existing event data to associate related events into a single incident object. Also, events that do not bear similarities on their face may also be combined or aggregated based upon event rate in a given network or sub-network. Thus varying the amount of event data sets destined for the analysis framework module).

Lyle merely teaches to associate related events. Lyle teaches: "The analysis framework 308 associates the event data with an event software object, as described more fully below, and stores data relating to the event in an event database 322. The analysis framework 308 also determines whether an event is associated with an existing event or group of related events, and associates related events into a single incident software object. Events that are not related to any other events are associated with a new incident object and may be later grouped with subsequently-received event data that is related to the same incident." [Lyle col.7, Line 61 to Col. 8 line 4]

Thus, Lyle does not describe varying the number of buckets according to the amount of traffic and number of flows into different buckets and examining statistics accumulated for a parameter and a corresponding threshold in the bucket.

Accordingly, since Lyle fails to describe all of the features of claim 63 arranged as in the claim, Lyle cannot anticipate claim 63.

Claims 64, 66, 68, 71, and 75

For the purposes of this appeal only claims 64, 66, 68, 71 and 75 stand or fall together. Claim 64 is representative of this group of claims.

Claim 64 further limits claim 63 and recites that: "varying varies the number of buckets so that the monitoring device is not vulnerable to DoS attacks against its own resources." This feature is not described by Lyle.

The examiner argues that: "As to claim 64, Lyle teaches the method of claim 63 wherein varying varies the number of buckets so that the monitoring device is not vulnerable to DoS attacks against its own resources (col 19, lines 37-45; the protocol disclosed by Lyle teaches a strong protection against denial of service attacks as well as other forms of attacks).

At Col. 19, lines 37-45, Lyle discloses: "In addition to this strong protection against this denial of service attacks the communication protocol described above protects the tracking systems from other types of attacks by requiring that the would be attacker both know the communication protocol and have the cryptographic hash function being used as part of the communication protocol in the tracking systems installed in the particular administrative domain."

However, as described by Lyle, it is not the event scheme that protects the tracking system from attacks but instead it is the: "In addition to this strong protection against this denial of service attacks the communication protocol described above protects the tracking systems from other types of attacks by requiring that the would be attacker both know the communication protocol and have the cryptographic hash function being used as part of the communication protocol." [Lyle, col. 19, lines 38-45]

Accordingly, since Lyle fails to describe all of the features of claim 64 arranged as in the claim, Lyle cannot anticipate claim 64.

Claims 65 and 72

For the purposes of this appeal only claims 65 and 72 stand or fall together. Claim 65 is representative of this group of claims.

Claim 65 further limits claim 63 and recites that varying the number of buckets includes comparing the number of buckets to a threshold number of buckets and determining whether the number of buckets should be divided into more buckets or combined into fewer buckets based on comparing the number of buckets to the threshold and as the number of buckets changes, the buckets have values derived from the buckets prior to the change.

The examiner contends that:

As to claim 65, Lyle teaches the method of claim 63 wherein varying the number of buckets comprises: comparing the number of buckets to a threshold number of buckets, determining whether the number of buckets should be divided into more buckets or combined into fewer buckets based on comparing the number of buckets to the threshold and as the number of buckets changes, the buckets have values derived from the buckets prior to the change (col 7, lines 43 to col 8, line 33; a statistics database is consulted including a threshold based upon incident rate to determine in part whether or not the event data set should be combined or split. Once a decision is made, variables within the event data set essentially remain the same).

Lyle does not describe that the number of buckets changes based on a comparison to a threshold. The examiner argues that: (col 7, lines 43 to col 8, line 33; a statistics database is consulted including a threshold based upon incident rate to determine in part whether or not the event data set should be combined or split. Once a decision is made, variables within the event data set essentially remain the same).” Lyle has no such teaching.

Lyle does not describe a threshold based on incident rate and does not determine whether event data should be combined or split based on a threshold. Rather, Lyle describes: “One of the tools used by analysis framework 308 in determining whether an event is associated with one or more other events is a statistics database 324. The statistics database 324 stores the average incident rate of each sub-network within the network served by the tracking system and a first-order variance of the average incident rate for all networks with an above-average incident rate. The baseline incident rate and the variance are used for all networks with an average or below-average incident rate.”

Lyle describes that: "The analysis framework 308 also connects to a policy database 326. The policy database 326 is used to store information concerning how certain types of events and incidents should be processed by the analysis framework, including the responsive action, if any, to be taken by the analysis framework. For example, for a particular type of attack or suspected attack the policy database 326 may indicate that the attack is to be logged but otherwise ignored." [Lyle, col. 8, lines 15-22]

Therefore according to Lyle, the incidence rate is used to process events. Lyle does not specifically describe that the incidence rate corresponds to a threshold number of buckets as in claim 1, but rather corresponds to the rate at which incidents occur in a network or sub-network. Therefore, Lyle does not describe determining whether the number of buckets should be divided into more buckets or combined into fewer buckets based on comparing the number of buckets to the threshold. Lyle also does not describe that as the number of buckets changes, the buckets have values derived from the buckets prior to the change.

Claims 67 and 74

For the purposes of this appeal only, claims 67 and 74 stand or fall together. Claim 67 is representative of this group of claims.

Claim 67 further limits claim 63 where comparing statistic values includes accumulating statistic values ... and comparing the values ... to thresholds that depend on the number of buckets. Lyle fails to suggest this feature. In Lyle the number of events is not based accumulating statistic values from the packets or comparing the values accumulated in the buckets. The examiner's reasoning that "(col 7, lines 3-20 and 43-67; sniffers are utilized in capturing packet content as well as data related to packets. Thereafter, the data requiring further analysis and/or evaluation is discerned and stored and placed into a queue for further scrutiny by the tracking system)." fails to address the claimed limitation.

**3. Claims 1-21, 50-62, 69, and 76-77 are
patentable over Lyle in view of Hsu et al.**

Claims 1 and 7

For the purposes of this appeal only, claims 1 and 7 stand or fall together. Claim 1 is representative of this group of claims.

Claim 1 calls a machine implemented method of monitoring traffic flow ... Claim 1 includes the features of producing statistics corresponding to a parameter of traffic flow to trace the source of an attack, ... mapping the traffic flow into a plurality of buckets by applying a hash function " $f(h)$ " to the parameter of the traffic flow to output an integer corresponding to one of the buckets, accumulating statistics from the packets; and comparing the number of buckets to a threshold. The claim also includes determining whether the number of buckets should be divided into more buckets or combined into fewer buckets based on comparing the number of buckets to the threshold.

The examiner contends with respect to these features that Lyle teaches: "Producing statistics corresponding to a parameter of traffic flow to trace the source of an attack (Fig 9, 908-310; col 2, lines 45-50; col 7, lines 3-12; sniffers are used in analyzing and evaluating traffic flows to scrutinize suspicious activity in an attempt to ascertain the source of an attack) ..."

For the reasons discussed above this feature is not taught by Lyle and Hsu does not cure the deficiencies in Lyle. Moreover, Lyle does not specifically suggest producing statistics corresponding to a parameter of traffic flow to trace the source of an attack, ... Accordingly to Lyle:

The sniffer module may also search for other information, clues, or signatures previously associated with attacks on the network being protected or other networks. For example, the sniffer module may identify all messages sent from one of a list of suspicious source addresses, or messages attempting to access a target system within the network or sub-network associated with the tracking system via a service known to be vulnerable, such as telnet, or messages containing strings present in messages associated with prior attacks.

Thus, Lyle discloses that the sniffer looks for strings of data present in packets or messages, and does not specifically suggest producing statistics corresponding to a parameter of traffic flow. Lyle further discusses that:

In one embodiment, statistical information from the statistics database is used to determine if the rate of certain types of messages, as described above exceeds a normal level. In one embodiment, the normal level or rate of certain types of message is programmed into the sniffer module as part of the configuration process and the sniffer module identifies as suspicious any series of data packets that exceed the rates established at the time of configuration

Lyle describes the statistical information as normal level or rate of certain types of messages. Thus, Lyle in no sense suggests much less describes producing statistics corresponding to a parameter of traffic flow, since Lyle merely examines patterns in the traffic flow not statistics on traffic flow.

Moreover, Lyle does not suggest mapping the traffic flow into a plurality of buckets. The examiner contends that: "Mapping the traffic flow into a plurality of buckets (col 7, lines 43-67; event data, which is defined as suspicious data is placed in a queue as a set corresponding to a single incident)." Event data however is neither mapped into buckets nor does the event data correspond to the traffic flow.

Lyle does not suggest: "Accumulating statistics from the packets and comparing the number of buckets to a threshold," whether at col 7, lines 32-42 and col 8, lines 6-14; or elsewhere. The examiner argues that: "many thresholds, such as incident rate, preconfigured criteria, timestamps, etc, are considered in determining the significance or importance of a possible attack." Whether that contention is correct or not, the contention does not address the claimed features namely accumulating statistics from the packets and comparing the number of buckets.

As for the feature of: "determining whether the number of buckets should be divided into more buckets or combined into fewer buckets based on comparing the number of buckets to the threshold." the examiner relies on col 7, line 43 to col 8, line 5; col 13, lines 42-50; and argues that: "once an event (a set of data corresponding to an attack) is placed in the queue, other event data is grouped or combined with existing event data to associate related events into a single

incident object. Also, events that do not bear similarities on their face may also be combined or aggregated based upon event rate in a given network or sub-network. Thus varying the amount of event data sets destined for the analysis framework module)."

The examiner's characterization does not address the feature of the claim, namely that the number of buckets are divided into more buckets or combined into fewer buckets based on comparing the number of buckets to the threshold. While indeed Lyle discloses a so-called single event object, that single incident object associates other event objects. Lyle does not describe or suggest that the number of objects are divided into more objects or combined into fewer objects. Rather, Lyle clearly discloses that the event objects are maintained in the log database. Thus, the single event object does not combine the objects but rather simply associates the objects for later analysis or retrieval.

Claim 1 also requires that the buckets are divided or combined base on a comparison to the threshold. Lyle does not suggest that the single incident object is formed based on comparing the number of buckets to a threshold such as the incidence rate.

The examiner acknowledges that "Lyle does not explicitly indicate the use of a hash function to output an integer corresponding to one of the buckets.", and thus relies on Hsu to teach "using a hash to output an integer corresponding to the location of a location of a unique bucket identifier (see fig 8, col 4, lines 26-38; col 5, lines 18-23)."

The examiner argues that:

It would have been obvious to one with ordinary skill in the art at the time the invention was made to combine the disclosure of Lyle with the hashing techniques in Hsu to make the system more efficient. Using the hashing technique, which utilizes addresses, will output the unique bucket identifier quickly. Because Lyle also uses addresses to relate event data to aggregate events into a single incident object, the use of Hsu's hashing technique would work seamlessly.

Applicant disagrees. Hsu describes that:

The routing information, comprising the source address 102 and the destination address 104, is then obtained 204 from the captured data packet 100. Byte size information of the captured data packet 100 is also preferably obtained 206. This information is then used to create or update 208 records stored in the memory of a conventional computer used to track the amount and size of data

traffic traveling into and out of specific nodes on the LAN, and the amount and size of data traffic traveling between two nodes on the LAN.

FIG. 3 shows a row 300 in an exemplary first table which contains information concerning the amount and size of data traveling into and out of a specific node on a LAN. Specifically, each row 300 from the table is indexed by a node address 302 and contains the following information: ...

Hsu is clearly directed to a table stored in memory and as such requires the use of a technique to distribute entries, e.g., a hash function. Lyle on the other hand is directed to an arrangement in which the incident objects are stored in a database, thus apparently being no such need for a distribution of entries. It would not be suggested to modify Lyle to hash the addresses of the entries, since Lyle uses a database. Indeed, Hsu does not describe "... mapping the traffic flow into a plurality of buckets by applying a hash function " $f(h)$ " to the parameter of the traffic flow to output an integer corresponding to one of the buckets," in any event.

Indeed, the examiner does recognize that: "Lyle does teach the use of hash functions in a unique way to efficiently communicate with the system (see col 19, lines 11-36), ..." However, appellant contends that since Lyle already recognized "hashing" it is not suggested to combine Lyle with Hsu to map the traffic flow into a plurality of buckets, which neither Lyle nor Hsu show, by applying the hash function to the parameter of the traffic flow to output an integer corresponding to one of the buckets, because that would not have advantaged the arrangement disclosed and suggested by Lyle.

Accordingly one of ordinary skill in the art would not be motivated to combine Lyle with Hsu and the combination, even if suggested does not teach all of the features of Applicant's claims. Accordingly, Hsu adds no further teachings to cure the deficiencies in Lyle and therefore the combination fails to suggest claim 1.

Claim 2

Claim 2 limits claim 1 and recites that the buckets are storage areas in memory. Lyle deals with a database and does not specifically discuss buckets as storage areas in memory. While events may reside in memory, temporarily they are ultimately logged into the database

taught by Lyle. Hsu, which does discuss memory, would not cure the deficiencies of Lyle, since it would change the principal of operation of Lyle and is therefore not suggested.

Claim 3

Claim 3 further limits claim 1 by reciting that as the number of buckets changes, the buckets have values derived from the buckets prior to the change. Claim 1 recites the feature of determining whether the number of buckets should be divided into more buckets or combined into fewer buckets based on comparing the number of buckets to the threshold. As discussed above, the combination of references do not suggest that the number of buckets changes based on a comparison to a threshold.

Appellant discusses in the specification (page 14, line 27) that:

The monitoring process 32 takes that bucket, e.g., B_i and divides that bucket B_i into some other number M of new buckets $B_{i1} - B_{iM}$. Each of the new buckets $B_{i1} - B_{iM}$ contains values appropriately derived from the original bucket B_i .

Appellant contends that Lyle does not teach that as the number of buckets changes, the buckets have values derived from the buckets prior to change whether at col 7, ln 59-67 or elsewhere, since the single incident object does not combine events but merely associates events, and thus the single incident object does not derived data from the events, but merely has the data from the events associated.

Claim 4

Appellant also describes (page 14, line 31) that:

Also, the hash function is extended to map to $N+M-1$ " $h \rightarrow N+M-1$ " values, rather than the original N values.
[Appellant's specification page 14, line 31 to page 15, line 2]

The examiner contends that: "As to claim 4, Lyle and Hsu teach the method claim of claim 1 wherein the hash function adapts to map to the new number of buckets, as the new

number of buckets changes (col 4, lines 26-38; the bucket identifier is unique and if a bucket is eliminated, so is its corresponding identifier. If, on the other hand, a bucket is added, a unique identifier is created).”

As discussed above the combination of Lyle and Hsu is not suggested and thus the concept of dividing buckets, and that the hash function adapts to map to the new number of buckets, as the new number of buckets changes, is not suggested.

Hsu at the cited passage teaches operation of a table that concerns information pertaining to traffic into or out of a node. However, Hsu does not disclose a record identifier at that passage. Hsu is devoid of any suggestion of a hash function that adapts to map to the new number of buckets, as the new number of buckets changes. Lyle also does not suggest a hash function that adapts based on changes in the number of buckets.

Claim 5

Claim 5 further limits claim 1 by comparing the value accumulated in the bucket to a threshold that depends on the number of buckets.

The examiner contends that

As to claim 5, Lyle teaches the method of claim 1, wherein comparing statistic values comprises accumulating statistic values from the packets and comparing the values accumulated in the buckets to thresholds that depend on the number of buckets, (col 7, lines 3-20 and 43-67; sniffers are utilized in capturing packet content as well as data related to packets. Thereafter, the data requiring further analysis and/or evaluation is discerned and stored and placed into a queue for further scrutiny by the tracking system).

Appellant disagrees. As discussed above Lyle does not suggest the features of claim 1 and thus does not suggest “comparing the value accumulated in the bucket to a threshold that depends on the number of buckets.” Lyle teaches to events and to associate events but in no sense does Lyle suggest to compare the value accumulated in the bucket to a threshold that depends on the number of buckets.

Claim 6

Claim 6 limits the method of claim 1 by reciting that the parameter is the count of how many packets a data collector or gateway examines. Lyle whether at col. 7, lines 3-20 or col. 7, lines 43-67 fails to suggest this feature. Lyle examines packets for strings or patterns. It is not seen where Lyle maintains any counts.

Claim 8

Claim 8 further limits claim 1 and requires that the hash function changes periodically in a randomly, secret manner so that packets are reassigned to different buckets.

As to claim 8, Lyle teaches the method of claim 1 wherein the hash function changes periodically in a randomly secret manner so that packets are reassigned to different buckets (Figs 11 A and B)

Lyle's discussion of the hash function and a random hash value pertain to the communication protocol, not to the features that the examiner relies on to suggest the features of claim 1. So although Lyle does disclose a hash function and a random number to use as a seed, Lyle does not disclose to apply that to reassignment of packets to different buckets in a randomly, secret manner.

Claim 9

Claim 9, further limits claim 1, and requires that the variable number of buckets dynamically adjusts as the amount of traffic and number of flows monitored so that the monitoring device is not vulnerable to a denial of service attack against its own resources. This feature is neither described nor suggest by Lyle whether at col. 19, lines 37-45 or elsewhere. While, the communications protocol disclosed by Lyle may be a strong protection against other forms of attack, Lyle does not disclose it as effective against denial of service attacks. Moreover, the communications protocol is not what the examiner uses in the rejection of claim 1 and thus the communication protocol feature of Lyle has no relevance to claim 9, since it is not seen that the event objects which the examiner does rely on provide the function or the features of claim 9.

Claim 10

Claim 10 depends from claim 1 and recites that the variable number of buckets efficiently identifies the source or sources of attack by breaking down traffic into different buckets and examining statistics accumulated for a parameter and a corresponding threshold in each bucket.

The examiner contends that:

As to claim 10, Lyle teaches the variable number of buckets efficiently identifies the source or sources of attack by breaking down traffic into different buckets and examining statistics accumulated for a parameter and a corresponding threshold in each bucket (col 8, lines 34-53; the event along with the policy assigned for that event is used in tracking the attack back to its origin, the incident object to which the event was designated would in fact identify the source of the attack).

Lyle fails to suggest to break traffic down into different buckets. Rather, Lyle tracks events and associates events into incident objects. Lyle has no teachings that suggest examining statistics accumulated for a parameter and a corresponding threshold in each bucket. Lyle teaches incident rates. However, these teachings are not a variable number of buckets that efficiently identifies the source or sources of attack by breaking down traffic into different buckets and examining statistics accumulated for a parameter and a corresponding threshold in each bucket.

Claim 11

Claim 11 further limits the method of claim 1 featuring that the traffic is monitored at multiple levels of granularity, from aggregate to individual flows. The examiner contends that Lyles teachings at col 7, lines 3 to col 8, line 53: "individual packets to events to incident objects are analyzed and evaluated at numerous times during processing of given information," correspond to monitoring traffic at multiple levels of granularity, from aggregate to individual flows. However Appellant notes that to the extent that events and incident objects correspond to multiple levels of granularity, those features of Lyle do not correspond to monitoring the traffic at multiple levels of granularity, from aggregate to individual flows.

Claim 12

Claim 12 further limits the method of claim 1 to where the method is applied to monitoring of TCP packet ratios and repressor traffic. The examiner argues that Lyle teaches this at "(col 7 line 59 to col 8, line 4; traffic from numerous types of networks including tcp/ip based networks is used and numerous values included in the statistics database are disclosed)."

Appellant disagrees. At no point in Lyle generally or at the cited passage does Lyle disclose "monitoring of TCP packet ratios and repressor traffic." Lyle monitors events, and discloses that: "When information related to an actual or suspected attack is received by the handoff receiver 302 or identified by the sniffer module 304, the relevant information is provided to an event manager module 306. The event manager 306 receives the suspicious data, referred to herein as "event" data, places it in a queue, and provides data to the analysis framework module 308 for processing, one event at a time, at predetermined intervals." Thus, Lyle neither suggests to process statistical information to determine the source of an attack nor the specific statistical information of claim 12, but rather only monitors events which correspond to an actual or suspected attack.

Claim 13

Claim 13 further limits the method of claim 1 by reciting that the threshold is a first threshold and the method includes comparing accumulated statistic values from the buckets to second threshold values to determine that an event is of significance.

The examiner contends that: "As to claim 13, Lyle teaches the method of claim 1 wherein further comprising comparing accumulated statistic values from the buckets to second threshold values to determine that an event is of significance (col 7, lines 32-42 and col 8, lines 6-14; many thresholds, such as incident rate, preconfigured criteria, timestamps, etc, are considered in determining the significance or importance of a possible attack)."

While Lyle teaches "baseline incident rate and a first order variance" this teaching does not suggest the first threshold as discussed above. However, claim 13 further requires that the second threshold is used with the accumulated statistical values from the buckets to determine that the event is of significance. While, the claimed second threshold is closer to the disclosed

baseline incident rate of Lyle, it still distinguishes over Lyle, since Lyle does not teach that the baseline incident rate is compared against accumulated statistic values from the buckets, but instead is used to compare the number of incidents of the event in a network. However, in no event does Lyle disclose the first and the second claimed thresholds, as required by claim 13.

Regarding Claims 14-21, 50-62 and 76-77 the examiner contended that they were: "essentially the computer program product and data collector for the above-mentioned method claims and are thus rejected under similar rationale."

Claims 14, 18, 19, 21, 53, 54, 57, 60, 61, 62, 77

For the purposes of this appeal only, claims 14, 18, 19, 21, 53, 54, 57, 60, 61, 62, 77 stand or fall together. Claim 14 is representative of this group of claims.

Claim 14 recites instructions to map, accumulate, compare and adjust, in an analogous manner as the corresponding features of claim 1. Claim 14 does not recite the feature of producing statistics, as in claim 1. Claim 14 distinguishes over the art since the cited references whether taken separately or in combination fail to suggest instructions to map the traffic flow into a plurality of buckets by applying a hash function to the parameter of the traffic flow to output an integer corresponding to one of the buckets." As argued above, event data however is neither mapped into buckets nor does the event data correspond to the traffic flow.

Claim 14 also distinguishes, since Lyle does not suggest instructions to accumulate statistics from the packets or instructions to compare the number of buckets to a threshold, whether at col 7, lines 32-42 and col 8, lines 6-14; or elsewhere. The examiner argues for claim 1 that: "many thresholds, such as incident rate, preconfigured criteria, timestamps, etc, are considered in determining the significance or importance of a possible attack." Whether that contention is correct or not, the contention does not address the claimed features namely accumulating statistics from the packets and comparing the number of buckets.

As for the feature of instructions to determine whether the number of buckets should be divided into more buckets or combined into fewer buckets based on comparing the number of buckets to the threshold.", this feature is not taught by Lyle whether at col 7, line 43 to col 8, line

5; col 13, lines 42-50 or elsewhere. The examiner's characterization and reliance on events does not address the feature that the number of buckets are divided into more buckets or combined into fewer buckets based on comparing the number of buckets to the threshold. While Lyle discloses a so-called single event object, that single incident object associates other event objects. Lyle does not describe or suggest that the number of objects are divided into more objects or combined into fewer objects. Rather, Lyle clearly discloses that the event objects are maintained in the log database. Thus, the single event object does not combine the objects but rather simply associates the objects for later analysis or retrieval.

Claim 14 also requires that the buckets are divided or combined base on a comparison to the threshold. Lyle does not suggest that the single incident object is formed based on comparing the number of buckets to a threshold such as the incidence rate.

The examiner acknowledges that "Lyle does not explicitly indicate the use of a hash function to output an integer corresponding to one of the buckets.", and thus relies on Hsu to teach "using a hash to output an integer corresponding to the location of a location of a unique bucket identifier (see fig 8, col 4, lines 26-38; col 5, lines 18-23)."

As argued above, Hsu is clearly directed to a table stored in memory and as such requires the use of a technique to distribute entries, e.g., a hash function. Lyle on the other hand is directed to an arrangement in which the incident objects are stored in a database, thus apparently being no such need for a distribution of entries. It would not be suggested to modify Lyle to hash the addresses of the entries, since Lyle uses the database.

Indeed, the examiner does recognize that: "Lyle does teach the use of hash functions in a unique way to efficiently communicate with the system (see col 19, lines 11-36), ..." However, appellant contends that since Lyle recognized "hashing" it is not suggested to combine Lyle with Hsu to map the traffic flow into a plurality of buckets by applying the hash function to the parameter of the traffic flow to output an integer corresponding to one of the buckets, because that would not have advantaged the arrangement disclosed and suggested by Lyle.

Accordingly one of ordinary skill in the art would not be motivated to combine Lyle with Hsu. Accordingly, Hsu adds no further teachings to cure the deficiencies in Lyle and therefore the combination fails to suggest claim 14.

Claims 15 and 50

For the purposes of this appeal only, claims 15 and 50 stand or fall together. Claim 15 is representative of this group of claims.

Claim 15 requires that based on the second threshold, the buckets are divided into more buckets or combined into fewer buckets. Lyle fails to disclose the claimed second threshold to divide or combine buckets.

Claims 16 and 51

For the purposes of this appeal only, claims 16 and 51 stand or fall together. Claim 16 is representative of this group of claims.

Claim 16 further limits claim 14 and recites instructions to divide the bucket into a different number of new buckets containing values derived from the original bucket. Lyle does not teach to divide buckets and because Lyle merely associates related events, would not inherently suggest dividing a bucket into a different number and derive values from the original bucket.

Claims 17 and 52

For the purposes of this appeal only, claims 17 and 52 stand or fall together. Claim 17 is representative of this group of claims.

Claim 17 further limits claim 14 to require the hash function adapt to map to the new number of buckets as the new number of buckets changes. Lyle does not teach the use of a hash to distribute entries as acknowledged by the examiner. The examiner contends that Hsu teaches this feature at (col. 9, lines 3-45).

Neither Hsu nor Lyle at the cited passages teach or have any suggestion of a hash function that adapts to map to the new number of buckets, as the new number of buckets changes.

Claims 20 and 55

For the purposes of this appeal only, claims 20 and 55 stand or fall together. Claim 20 is representative of this group of claims.

Claim 55 further limits claim 21 to a hash function that changes periodically in a randomly secret manner so that packets are reassigned to different buckets. Lyle, as admitted by the examiner, fails to suggest the claimed hash function. Neither Lyle nor Hsu would have any use for a secret hash function for mapping, since Lyle merely uses the hash function as part of the communication protocol whereas Hsu uses the hash to distribute entries in a table, but does not suggest any need for secrecy in how data records are distributed in the table, apparently since Hsu collects flows from devices on the network apparently as part of a tool to, e.g., where: "... the network can be analyzed and possibly redesigned for improved transmission of data packets across the network."

Claim 56

Claim 56 further limits claim 21 by comparing the value accumulated in the bucket to a threshold that depends on the number of buckets. This claim is allowable for analogous reasons given in claim 5.

Claim 58

Claim 58 further limits claim 21 reciting that the variable number of buckets dynamically adjusts the amount of traffic and number of flows monitored, so that the data collector is not vulnerable to a denial of service attack against its own resources. Lyle fails to suggest any mechanism to protect the tracking system from a denial of service attack against its own resources. Variable number of buckets dynamically prevents such an exploit.

Claim 59

Claim 59 further limits claim 21 to a data collector that uses the variable number of buckets to efficiently identify the source or sources of an attack by breaking down traffic into different buckets and examining statistics accumulated for a parameter and a corresponding threshold in each bucket. Lyle does not suggest this feature. Lyle uses events and "Messages

associated with an attack may be tracked back to identify a point of attack at which messages associated with the attack are entering a network." However, Lyle does not use buckets to determine the source of an attack, but rather, e.g., "'sniffer' module comprised of one or more 'sniffers', described more fully below, continuously scans the data being received at various ports of various network devices. The sniffers search for data indicating an actual or suspected attack, as described more fully below, and provide information concerning suspicious data to other modules within the tracking system, as described more fully below." In contrast, Claim 59 requires that the statistics accumulated for a parameter and a corresponding threshold in each bucket are used to identify the source of an attack.

Claims 69 and 76

For the purposes of this appeal only, claims 69 and 76 stand or fall together. Claim 69 is representative of this group of claims.

Claim 69 limits the method of claim 63 wherein the buckets are storage areas in a memory space of the monitor device and mapping the traffic flow into a plurality of buckets comprises applying a hash function " $f(h)$ " to the parameter of the traffic flow to output an integer corresponding to one of the buckets.

Lyle fails to suggest applying a hash function, as admitted by the examiner and the modification of Lyle with Hsu is not suggested nor provides "mapping the traffic flow into a plurality of buckets comprises applying a hash function " $f(h)$ " to the parameter of the traffic flow to output an integer corresponding to one of the buckets."

Conclusion


Appellant submits, therefore, that Claims 1-30 and 32 are allowable over the cited art. Therefore, the Examiner erred in rejecting Appellant's claims and should be reversed.

Respectfully submitted,

Date: _____

11/6/06

Denis G. Maloney
Reg. No. 29,670



Fish & Richardson P.C.
225 Franklin Street
Boston, MA 02110-2804
Telephone: (617) 542-5070
Facsimile: (617) 542-8906

Appendix of Claims

1. A machine implemented method of monitoring traffic flow in a monitoring device disposed to receive network traffic packets comprises:

producing statistics corresponding to a parameter of traffic flow to trace the source of an attack, with producing further comprising:

mapping the traffic flow into a plurality of buckets by applying a hash function " $f(h)$ " to the parameter of the traffic flow to output an integer corresponding to one of the buckets;

accumulating statistics from the packets; and

comparing the number of buckets to a threshold; and

determining whether the number of buckets should be divided into more buckets or combined into fewer buckets based on comparing the number of buckets to the threshold.

2. The method of claim 1 wherein the buckets are storage areas in a memory space of the monitor device.

3. The method of claim 1 wherein as the number of buckets changes, the buckets have values derived from the buckets prior to the change.

4. The method of claim 1 wherein the hash function adapts to map to the new number of buckets, as the new number of buckets changes.

5. The method of claim 1 wherein comparing statistic values comprises:
comparing the value accumulated in the bucket to a threshold that depends on the number of buckets.

6. The method of claim 1 wherein the parameter is the count of how many packets a data collector or gateway examines.

7. The method of claim 1 wherein as a value of a parameter for one bucket approaches a threshold, the monitoring device raises an alarm.

8. The method of claim 1 wherein the hash function changes periodically in a randomly secret manner so that packets are reassigned to different buckets.

9. The method of claim 1 wherein the variable number of buckets dynamically adjusts the amount of traffic and number of flows monitored, so that the monitoring device is not vulnerable to a denial of service attack against its own resources.

10. The method of claim 1 wherein the variable number of buckets efficiently identifies the source or sources of attack by breaking down traffic into different buckets and examining statistics accumulated for a parameter and a corresponding threshold in each bucket.

11. The method of claim 1 wherein the traffic is monitored at multiple levels of granularity, from aggregate to individual flows.

12. The method of claim 1 wherein the method is applied to monitoring of TCP packet ratios and repressor traffic.

13. The method of claim 1 wherein the threshold is a first threshold and the method further comprises:

comparing accumulated statistic values from the buckets to second threshold values to determine that an event is of significance.

14. A computer program product residing on a computer readable for monitoring network traffic flow in a network comprises instructions for causing a computer to:

map traffic flow into a plurality of buckets by applying a hash function "f(h)" to a parameter of the traffic flow to output an integer corresponding to one of the buckets;
accumulate statistics from the packets; and
compare the accumulated statistic values from the buckets to configured threshold values corresponding to the number of buckets to determine that an event is of significance; and
adjust the number of buckets as the number of buckets approaches a second threshold.

15. The computer program product of claim 14 wherein based on the second threshold, the buckets are divided into more buckets or combined into fewer buckets

16. The computer program product of claim 14 wherein instructions to monitor further comprise instructions to
divide the bucket into a different number of new buckets containing values derived from the original bucket.

17. The computer program product of claim 14 wherein the hash function adapts to map to the new number of buckets as the new number of buckets changes.

18. The computer program product of claim 14 wherein the parameter is the count of how many packets a data collector or gateway examines.

19. The computer program product of claim 14 wherein the buckets are storage areas in the memory space of the monitor device.

20. The computer program product of claim 14 wherein the hash function changes periodically in a randomly secret manner so that packets are reassigned to different buckets.

21. A data collector to collect statistical information about network flows comprises:
a computer readable medium;
a computing device that executes a computer program product stored on the computer readable medium comprising instructions to cause the computing device to:
map traffic flow into a plurality of buckets by applying a hash function " $f(h)$ " to the parameter of the traffic flow to output an integer corresponding to one of the buckets;
accumulate statistics from the packets; and
compare the accumulated statistic values from the buckets to configured threshold values corresponding to the number of buckets to determine that an event is of significance; and
adjust the number of buckets as the number of buckets approaches a second threshold.

Claims 22-49 are canceled.

50. The data collector of claim 21 wherein based on the second threshold, the buckets are divided into more buckets or combined into fewer buckets

51. The data collector of claim 21 wherein instructions to monitor further comprise instructions to
divide the bucket into a different number of new buckets containing values derived from the original bucket.

52. The data collector of claim 21 wherein the hash function adapts to map to the new number of buckets as the new number of buckets changes.

53. The data collector of claim 21 wherein the parameter is the count of how many packets the data collector examines.

54. The data collector of claim 21 wherein the buckets are storage areas in the memory space of the monitor device.

55. The data collector of claim 21 wherein the hash function changes periodically in a randomly secret manner so that packets are reassigned to different buckets.

56. The data collector of claim 21 wherein instructions to compare statistic values comprises instructions to:

compare the value accumulated in the bucket to a threshold that depends on the number of buckets.

57. The data collector of claim 21 wherein as a value of a parameter for one bucket approaches a threshold, the monitoring device raises an alarm.

58. The data collector of claim 21 wherein the variable number of buckets dynamically adjusts the amount of traffic and number of flows monitored, so that the data collector is not vulnerable to a denial of service attack against its own resources.

59. The data collector of claim 21 wherein the variable number of buckets efficiently identifies the source or sources of attack by breaking down traffic into different buckets and examining statistics accumulated for a parameter and a corresponding threshold in each bucket.

60. The data collector of claim 21 wherein the traffic is monitored at multiple levels of granularity, from aggregate to individual flows.

61. The data collector of claim 21 wherein the traffic is applied to monitoring of TCP packet ratios and repressor traffic.

62. The data collector of claim 21 wherein the threshold is a first threshold and the computer program further comprises instructions to:

compare accumulated statistic values from the buckets to second threshold values to determine that an event is of significance.

63. A method of monitoring traffic flow in a monitor device disposed to receive network packets, the method comprises:

producing statistics corresponding to a parameter of the traffic flow to trace a source of an attack, with producing further comprising:

mapping the traffic flow into a plurality of buckets;

varying the number of buckets according to the amount of traffic and number of flows to breakdown traffic flow into different buckets; and

analyzing statistics accumulated for a parameter and a corresponding threshold in the bucket to identify the source of the attack.

64. The method of claim 63 wherein varying varies the number of buckets so that the monitoring device is not vulnerable to DoS attacks against its own resources.

65. The method of claim 63 wherein varying the number of buckets comprises:

comparing the number of buckets to a threshold number of buckets;

determining whether the number of buckets should be divided into more buckets or combined into fewer buckets based on comparing the number of buckets to the threshold and as the number of buckets changes, the buckets have values derived from the buckets prior to the change.

66. The method of claim 63 wherein further comprising:

comparing accumulated statistic values from the buckets to second threshold values to determine that an event is of significance.

67. The method of claim 63 wherein comparing statistic values comprises:
accumulating statistic values from the packets; and
comparing the values accumulated in the buckets to thresholds that depend on the number of buckets.

68. The method of claim 63 wherein the variable number of buckets dynamically adjusts the amount of traffic and number of flows monitored, so that the monitoring device is not vulnerable to a denial of service attack against its own resources.

69. The method of claim 63 wherein the buckets are storage areas in a memory space of the monitor device and mapping the traffic flow into a plurality of buckets comprises:
applying a hash function " $f(h)$ " to the parameter of the traffic flow to output an integer corresponding to one of the buckets.

70. A computer program product residing on a computer readable medium for monitoring traffic flow in a monitor device disposed to receive network packets, the computer program product comprises instructions for causing the device to:
produce statistics corresponding to a parameter of the traffic flow to trace a source of an attack, with producing further comprising:
map the traffic flow into a plurality of buckets;
vary the number of buckets according to the amount of traffic and number of flows to
breakdown the traffic flow into different buckets; and
analyze statistics accumulated for a parameter and a corresponding threshold in the bucket to identify a source of the attack.

71. The computer program product of claim 70 wherein instructions to vary, vary the number of buckets so that the monitoring device is not vulnerable to DoS attacks against its own resources.

72. The computer program product of claim 70 wherein instructions to vary comprises instructions to:

compare the number of buckets to a threshold number of buckets;

determine whether the number of buckets should be divided into more buckets or combined into fewer buckets based on comparing the number of buckets to the threshold and as the number of buckets changes, the buckets have values derived from the buckets prior to the change.

73. The computer program product of claim 70 further comprising instructions to:

compare accumulated statistic values from the buckets to second threshold values to determine that an event is of significance.

74. The computer program product of claim 70 wherein instructions to compare statistic values comprises instructions to:

accumulate statistic values from the packets; and

compare the values accumulated in the buckets to thresholds that depend on the number of buckets.

75. The computer program product of claim 70 wherein the variable number of buckets dynamically adjusts the amount of traffic and number of flows monitored, so that the monitoring device is not vulnerable to a denial of service attack against its own resources.

76. The computer program product of claim 70 wherein the buckets are storage areas in a memory space of the monitor device and instructions to map the traffic flow into a plurality of buckets comprises instructions to:

apply a hash function "f(h)" to the parameter of the traffic flow to output an integer corresponding to one of the buckets.

77. The data collector of claim 21 further comprising:
a port to link the data collector to a central control center.

Applicant : Thomas Michael Gil et al.
Serial No. : 09/931,223
Filed : August 16, 2001
Page : 39 of 40

Attorney's Docket No.: 12221-007001

Evidence Appendix

None

Applicant : Thomas Michael Gil et al.
Serial No. : 09/931,223
Filed : August 16, 2001
Page : 40 of 40

Attorney's Docket No.: 12221-007001

Related Proceedings Appendix

None